920476-904962

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In the application of	:	Rory Smith
Serial No.	:	09/975,841
Filed	:	October 12, 2001
For	:	Bandwidth allocation in a synchronous transmission network for packet oriented signals
Examiner	:	Nguyen Hoang Ngo
Art Unit	:	2616
Customer number	:	23644

BRIEF ON APPEAL

Honorable Director of Patents and Trademarks
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

This appeal is from the Examiner's final rejection of May 3, 2006 and the advisory action of July 13, 2006 in which all pending claims, that is claims 1-28 and 30-36 were rejected or objected to. An appropriate response was e-filed on June 27, 2006, and a timely Notice of Appeal was mailed on August 3, 2006 with the required fee of $500, following an Advisory Action of July 13, 2006.

The required $500 fee pursuant to 37 C.F.R. § 41.20(b)(2) may be deducted from Deposit Account No. 12-0913

(i) Real Party in Interest

This application is assigned to Nortel Networks Limited.

(ii) Related Appeals and Interferences

There are no related appeals or interferences or judicial proceedings.

(iii) Status of Claims

This application was filed with claims 1 through 36. Claim 29 was deleted, and Claims 1-28 and 30-36 have been finally rejected by the Examiner. The rejection of these claims in the final Office Action of May 3, 2006 is appealed. Claims 1-28 and 30-36 as amended during the prosecution of the application are set forth in the Claims Appendix.

(iv) Status of Amendments

No amendments were made in the response of June 27, 2006. Only arguments were presented.

(v) Summary of Invention

The present invention relates generally to communication networks, and more specifically to efficiently using synchronous networks for passing packet oriented signals from one node of a packet oriented network to another, using buffer-to-buffer flow control.

As is explained in the specification at page 12 onwards, buffer-to-buffer flow control regulates traffic along a link between a transmitter port and a receiver port by controlling the rate at which the transmitter can send data to the receiver. A

transmitter is able to transmit a frame along a link only if the receiver has indicated it can accept the frame. Buffering may be required for any one of a number of reasons. For example, the client signal may need to be buffered to mitigate the effects of any congestion along its route/at its destination. In the present case, buffering can be used to manage handing-off between the different client signal data rates and synchronous network payload rates. The distinctive feature of preserving the client signal buffer-to-buffer flow control when processing to map the signal to the synchronous network can help make the synchronous network transparent to the client signal, meaning that client signal nodes at each side of the synchronous link can appear on the same client network. This helps simplify management of the client network 100, 110.

Independent claim 1:

This claim specifies a method of mapping a packet orientated client signal to a synchronous network payload, the method including the steps of:

receiving said client signal (F1);

processing (210, 212) said client signal to a form suitable for mapping to said payload which preserves a buffer-to-buffer flow control mechanism (300, 310a) of the client signal, wherein said step of processing reduces (212) the bandwidth of the client signal while maintaining the integrity of a payload of the client signal;

and mapping (218) said processed signal to said synchronous network payload.

Reference numerals have been inserted based on figs 3A and 5, which will be discussed briefly now. Figure 3A shows an overview of networks having a number of nodes, and in particular communications paths between node A of a packet oriented network 100, passing through synchronous network 120, to reach node Z of another packet oriented network 110. The claimed steps would take place at the interface 114 to the synchronous network, at node C or node N1, for communication from A to Z. The buffer to buffer flow control of the packet oriented client signal, which is

preserved through the synchronous network, is from a buffer at node C, to a buffer at node X.

Fig 5 describes an example of the operation at node C or at node N1. The client signal F1 is received at the node as indicated by the arrow leading to first box 300. This box and box 310a are concerned with the buffer to buffer flow control mechanism, for controlling flow between buffers C and X. The claimed step of processing while preserving the flow control mechanism is shown by steps 210 and 212. The claimed step of mapping to the synchronous payload is shown by box 218 which follows box 212. This leads to transmission over the synchronous network at box 220. The claim feature of the processing preserving the buffer to buffer flow control mechanism is shown for example the processing steps 210 and 212 in Fig 3B being dependent on the flow control mechanism steps 300 and 310a.

This flow control mechanism is described in more detail at page 17 line 25 onwards which indicates that "The receiving port controls the transmission of frames by giving permission to the sending port to send one or more frame to that particular receiving port. That permission is called a credit. The number of frames that may be sent is referred to as the available credit. Flow control is provided by both ports on the link exchanging the number of frames they may receive at any time from each other to determine their respective buffer credit values during a "log on" phase."

At page 18 line 12, it is explained that "Each port keeps track of the buffer credit count, which is initialised to zero. For each frame transmitted, the credit count is incremented by one, and for each R_RDY primitive signal received from the other port the credit count is decreased by one. Transmission of a R_RDY primitive signal indicates a port has processed a frame, freed a receive buffer, and is ready for one more. If the buffer credit count reaches the buffer credit value at the initiator port, no more frames can be sent by the initiator...".

An aim of the claimed features is to enable more efficient transmission of packet oriented protocols which can use a buffer-to-buffer flow control mechanism, such as Fibre Channel and ESCON, through a synchronous network. Such a flow control mechanism helps to ensure that the buffers of receiving and transmitting ports along a particular communications path do not overflow, which could cause data to be lost and/or retransmitted. Notably, preserving the client flow control mechanism through the synchronous network helps avoid the need for complex protocol specific functions. In particular the step of reducing the bandwidth while maintaining the payload integrity before the mapping can avoid the need for complex flow control or buffer credit management functionality, or any intermediate mapping to other OSI layer -2 protocols such as ATM. Notably the buffers of the buffer-to-buffer mechanism referred to in the claim are buffers controlled by the mechanism of the client signal.

<u>Independent claim 16</u>:

This claim has corresponding distinctive features and so the discussion of claim 1 applies here. It specifies a method of mapping a packet oriented client signal that uses a buffer-to-buffer flow control mechanism to a synchronous transmission network payload, the method comprising the steps of:

processing said client signal (212) to remove at least one ordered set provided according to a protocol of said client signal to form a second signal;

storing the second signal in an ingress buffer (214); and

mapping (216a, 216b, 218) the second signal to said synchronous payload,

wherein said steps of processing said client signal and mapping said second signal preserves the buffer-to-buffer flow control mechanism of the client signal and maintains the integrity of the payload of the client signal.

Reference numerals have been added based on Fig 4B, which shows operations at node N1. This shows a similar process to that described above with relation to Fig 3A, and so need not be described again here.

Independent claim 20:

Again, this claim has corresponding distinctive features and so the discussion of claim 1 applies here. It specifies a method of restoring a packet oriented client signal from at least one synchronous network payload, the method comprising the steps of:

receiving (222) said synchronous payload;

de-mapping (222, 224a, 224b) said signal from said synchronous payload;

storing (226) said signal in an egress buffer; and

processing (228, 230) said signal to add at least one ordered set provided according to a protocol of said packet orientated client signal, wherein said method of restoring the client signal maintains the integrity of the payload of said packet oriented client signal and preserves a buffer-to-buffer flow control mechanism of said client signal.

Reference numerals have been added based on the right hand side of fig 4B. This shows the steps at node N4, where the path leaves the synchronous network to reach the packet oriented network 110. This shows a corresponding process to that of Fig 3A, in the reverse order, and so need not be described again in detail.

Independent claim 24:

Again, this claim has corresponding distinctive features and so the discussion of claim 1 applies here. It specifies apparatus adapted to perform steps in a method of mapping a client signal comprising a packet oriented client signal (F1) that uses a buffer-to-buffer flow control mechanism to a synchronous transmission network payload, the apparatus comprising:

a processor (210, 212) for processing said client signal to remove at least one ordered set provided according to a protocol of said client signal to form a second signal;

a buffer (214) for storing the processed client signal in an ingress buffer; and

a mapper (216a, 216b, 218) for mapping the processed client signal to said synchronous payload,
wherein said apparatus preserves the buffer-to-buffer flow control mechanism of the client signal and maintains the integrity of the payload of the client signal.

Reference numerals have been added based on the left hand side of fig 4B. This shows a similar process to that described above with relation to Fig 3A, and so need not be described again here.

Independent claim 30:
Again, this claim has corresponding distinctive features and so the discussion of claim 1 applies here. It specifies a method of load balancing traffic comprising a packet orientated client signal across a synchronous network (120), wherein said traffic comprises at least one synchronous network payload comprising a packet oriented client signal which is controlled by a buffer-to-buffer flow control mechanism (300, 310a), the signal having been mapped to a synchronous network payload, using a method including the steps of: receiving said client signal (F1); processing (210, 212) said client signal to a form suitable for mapping to said payload which preserves a buffer-to-buffer flow control mechanism of the client signal, wherein said step of processing reduces the bandwidth of the client signal while maintaining the integrity of a payload of the client signal; and mapping (218) said processed signal to said synchronous network payload, wherein said method of load balancing comprises the steps of:

pre-allocating (by node N1, fig 3A, as described at page 33) an initial bandwidth of said synchronous network payload according to a predetermined condition , wherein said payload comprises a plurality of virtually concatenated virtual containers;

diversely routing (by node N1, fig 3A, as described at page 33) said synchronous network payload over said synchronous network; and

in the event of a change in a condition of the network, modifiying the allocated bandwidth (by node N1, fig 3A, as described at page 33).

Reference numerals have been added based on figs 3A and 5 and so these features need not be described again here.

Independent claim 36:

Again, this claim has corresponding distinctive features and so the discussion of claim 1 applies here. It specifies a method of allocating bandwidth in a synchronous digital network for a packet oriented signal (F1) having buffer-to-buffer flow control (300, 310a), the method comprising the steps of:

received said packet oriented signal;

processing (210, 212) said packet oriented signal to a processed signal having a form suitable for mapping to a synchronous network payload, wherein the processing preserves a buffer-to-buffer flow control mechanism (300, 310a) of said packet oriented signal, wherein said step of processing removes redundant information from the packet oriented signal while maintaining the integrity of a payload of the packet oriented signal;

and mapping (218) said processed signal to a said synchronous network payload having a bandwidth determined according to the bandwidth of said processed signal.

Reference numerals have been added based on figs 3A and 5 and so these features need not be described again here.

(vi) Grounds of Rejection to be Review on Appeal

There are two grounds of rejection:

(1). Claims 1-6, 8, 11, 12, 16-20, 22, 23,24-28 and 36 have been rejected under 35 U.S.C. §102(e) for anticipation by Jordan (US 6,934,301).

(2). The remaining claims are rejected for obviousness over Jordan and various other references. It seems that all rejections are dependent on the Examiner's interpretation of Jordan as showing a buffer- to- buffer flow control. If it does not show this feature, then it cannot show the claim feature of preserving this mechanism when processing for mapping to a payload of a synchronous network.

(i) Claim 7 has been rejected under 35 U.S.C. §103 as being unpatentable over Jordan in view of Rauch (U.S. 6,243,510).

(ii) Claims 9 and 15 have been rejected under 35 U.S.C. §103 as being unpatentable over Azizoglu (U.S. 6,430,201).

(iii) Claims 10 and 21 have been rejected under 35 U.S.C. §103 as being unpatentable over Jordan alone.

(iv) Claim 13 has been rejected under 35 U.S.C. §103 as being unpatentable over Jordan in view of Gerszberg (U.S. 6,452,923).

(v) Claim 14 has been rejected under 35 U.S.C. §103 as being unpatentable over Jordan in view of Goodman (U.S. 6,636,529).

(vi) Claims 30-35 have been rejected under 35 U.S.C. §103 as being unpatentable over Jordan in view of Heuer (U.S. 6,842,455).

(viii) Argument

Ground 1

In summary, Jordan only shows a conventional buffer for absorbing rate changes at an interface from the packet network to the synchronous network. It does not show a buffer-to-buffer flow control. Even if it did, it does not show the claim feature of processing the signal to a form suitable for mapping to said synchronous payload, to preserve the flow control mechanism. This cannot be shown, because any buffer-to-buffer flow control in Jordan can only occur before any mapping to the synchronous network and so there is no need for, and therefore no implied disclosure of, preserving the flow control mechanism.

These issues will now be explained in more detail by reference to earlier arguments. In the response of March 30, 2006, concerning claim 1, it was explained that Jordan is concerned generally with converting a conventional data packet received from a 1 Gb Ethernet network to a conventional data packet suitable for transmitting on a conventional standard bandwidth synchronous SONET network such as an OCnc (n=1, 3, 12) payload network with no loss of data content. This is achieved by receiving a series of data packet bursts from a broadband network with idle bytes interposed between the bursts; removing the idle bytes to reduce a transmitted bit stream, framing the packets in accordance with a conventional protocol such as a General Frame Protocol (GFP) or Packet Over Sonet protocol (POS), and providing the framed data packets to the synchronous payload network.

The cited passage of Jordan is only a conventional buffer overflow control. It is not concerned with a buffer-buffer mechanism, since there is only one buffer, and no disclosure of any buffer at the sending side. Furthermore, it is not concerned with a flow control mechanism of the client signal, since the client signal is not involved in

the buffer overflow. This converter buffer of Jordan is not visible to or controlled by the client signal. Instead this passage of Jordan is only concerned with a mechanism for stopping sending from a source if a buffer at the converter becomes filled.

The cited passage states that:

> "Since the gate keeping signal provided by control logic block 120 corresponds to the 1 Gb clock rate of the Ethernet network and is therefore faster than the signal provided by OCnc payload clock 130, the rate at which data is written to buffer 116 can be faster than the rate at which data is read from buffer 116. Thus, and in order to prevent a buffer overflow, the values of read pointer 122 and write pointer 124 are provided to register 126 to regulate enabling the gate keeping signal of control logic block 120. A conventional pulse command (not shown) can also be returned to the Ethernet network instruction it to stop sending data. Thus, buffer 116 absorbs the differential between the write rate from the Ethernet network 50 and the read rate to the OCnc payload network 60."

Hence there is no disclosure in Jordan of the claim feature of a buffer-buffer flow control mechanism of the client signal. Hence there cannot be any disclosure of the claim feature of:

"processing said client signal to a form suitable for mapping to said payload which preserves a buffer-to-buffer flow control mechanism of the client signal..."

Hence claim 1 is not anticipated. As Jordan relates only to Ethernet which does not have such a flow control mechanism, and does not relate to a protocol such as Fiberchannel or ESCON, which do have a buffer-buffer flow control mechanism, there is nothing in Jordan leading a skilled person to adapt it to reach the claimed invention.

In the final Office Action of March 5, 2006, the Examiner did not accept these explanations and tried to argue that "buffer-to-buffer flow control" can be interpreted as a "buffer overflow control of different Ethernet signals. " In response, it was explained in detail that buffer-to-buffer must mean from one buffer to another buffer. It cannot logically mean one buffer to the same buffer even if the same buffer is used by two different signals. If the Examiner regarded that "buffer-to-buffer flow control" can encompass multiple buffers in parallel each having their own conventional single buffer overflow control, this might be called a "buffer-by-buffer flow control". But this cannot be a logical interpretation of a "buffer-to-buffer flow control" because the word "-to-" has a distinct meaning from "-by-" .That a skilled person could not logically interpret "buffer-to-buffer flow control" as "buffer-by-buffer flow control" is confirmed in this case by the fact that the claim specifies that the flow control is "of the client signal". It would have to say "of the signals" plural, for buffer-by-buffer flow control to be a possible interpretation, and even then there would have to be some reason for ignoring the usual meaning of "-to-".

In response, in the Advisory Action of July 13, 2006, the Examiner essentially reiterates earlier arguments. In particular, the Examiner seems to try to argue that a buffer-to-buffer flow control can encompass different Ethernet signals passing through the same buffer.

Notably, the Examiner has not commented on or tried to rebut the above arguments that this claim term must mean one buffer to a next buffer, and cannot logically mean one buffer to the same buffer even if the same buffer is used by two different signals. The Examiner goes on in the last two sentences of item 2 to try to argue that Jordan does show control of one buffer affecting flow of another buffer in the system. This presumably refers to another buffer earlier in the Ethernet network. This implies the Examiner is persuaded by the logic that buffer-to-buffer must mean one buffer to a different buffer. But Jordan does not mention another buffer, so this is speculation, based only on the passage cited above that "A conventional pulse command (not

shown) can also be returned to the Ethernet network instruction it to stop sending data." There might or might not be another buffer earlier, Jordan is silent about this.

In any case, the claim also specifies that the signal is processed to a form suitable for mapping to said synchronous payload, to preserve the flow control mechanism. Even if Jordan were to be interpreted as showing a buffer-to-buffer flow control mechanism, it is between buffers before the converter to the synchronous network. There is no suggestion of buffers after the synchronous network. Thus any implied flow control must end before the converter, and thus end before any mapping. Hence there is no disclosure of any mapping to the synchronous network which preserves the flow control. Hence claim 1 cannot be anticipated by Jordan.

Regarding obviousness, as there is no suggestion in Jordan leading towards flow control through a synchronous network, there cannot be any suggestion of the claim feature of a mapping to the synchronous network which preserves the flow control, and so there cannot be any incentive to lead a skilled person to alter Jordan to reach the claimed invention. Hence claim 1 cannot be obvious over Jordan. None of the other documents are any more relevant, and so it cannot be obvious over any combination of documents.

<u>Ground 2</u>

All the other claims are dependent or have corresponding distinctive features and so are all allowable for the same reasons.

Reversal of the Examiner's rejections is respectfully requested.

October 2, 2006

Respectfully submitted,

William M. Lee, Jr.
Registration No. 26,935
Barnes & Thornburg LLP
P.O. Box 2786
Chicago, Illinois 60690-2786
(312) 214-4800
(312) 759-5646 (fax)

Claims Appendix

1. A method of mapping a packet orientated client signal to a synchronous network payload, the method including the steps of:

receiving said client signal;

processing said client signal to a form suitable for mapping to said payload which preserves a buffer-to-buffer flow control mechanism of the client signal, wherein said step of processing reduces the bandwidth of the client signal while maintaining the integrity of a payload of the client signal;

and mapping said processed signal to said synchronous network payload.

2. A method as claimed in claim 1, wherein the bandwidth is reduced by removing redundant information from said client signal.

3. A method as claimed in claim 1, wherein the bandwidth is reduced by removing idles from said client signal.

4. A method as claimed in claim 1, wherein the bandwidth is reduced by removing at least one primitive sequence which forms part of a series of repeated primitive sequences in said client signal.

5. A method as claimed in claim 1, wherein in the step of preserving the buffer-to-buffer flow control mechanism of the client signal, said buffer-to-buffer flow control mechanism is provided according to a Fibre Channel protocol class of service.

6. A method as claimed in claim 1, wherein in the step of preserving the buffer-to-buffer flow control mechanism of the client signal, said buffer-to-buffer flow control mechanism is provided according to an ESCON protocol class of service.

7. A method as claimed in claim 5 , wherein said packet orientated client signal is provided according to a higher level protocol supported by said Fibre Channel protocol and which has a buffer-to-buffer flow control mechanism provided according to a Fibre Channel protocol class of service.

8. A method as claimed in claim 1, wherein the synchronous payload is taken from the group consisting of: one or more SONET virtual container payloads, one or more SDH virtual container payloads; two or more virtually concatenated SONET

virtual container payloads; two or more virtually concatenated SDH virtual container payloads; two or more contiguously concatenated SONET virtual container payloads; two or more contiguously concatenated SDH virtual container payloads.

9. A method as claimed in claim 1, wherein said step of processing the client signal further includes a step of removing line encoding.

10. A method as claimed in claim 1, further including the step of padding said processed client signal so that said processed client signal is appropriately padded to fill a predetermined synchronous payload bandwidth.

11. A method as claimed in claim 1, wherein the bandwidth of the synchronous payload is allocated by a network management system.

12. A method as claimed in claim 1, wherein the bandwidth of the synchronous payload is allocated by an apparatus implementing the method of mapping.

13. A method of mapping as claimed in claim 1, wherein the synchronous payload bandwidth is modified in response to customer bandwidth demands increasing/decreasing.

14. A method of mapping as claimed in claim 1, wherein the synchronous payload bandwidth is modified in response to changes in data throughput as distance between the end data packet nodes changes.

15. A method as claimed in claim 1, wherein a plurality of clients signals are multiplexed together to share said synchronous payload.

16. A method of mapping a packet oriented client signal that uses a buffer-to-buffer flow control mechanism to a synchronous transmission network payload, the method comprising the steps of:

processing said client signal to remove at least one ordered set provided according to a protocol of said client signal to form a second signal;

storing the second signal in an ingress buffer; and

mapping the second signal to said synchronous payload,

wherein said steps of processing said client signal and mapping said second signal preserves the buffer-to-buffer flow control mechanism of the client signal and maintains the integrity of the payload of the client signal.

17. A method as claimed in claim 16, wherein said ordered set provides redundant data in said client signal.

18. A method as claimed in claim 16, wherein said ordered set provides redundant data comprising at least one client signal idle.

19. A method as claimed in claim 16, wherein said ordered set provides redundant data comprising at least one client signal primitive sequence which is repeated in a series of client signal primitive sequences .

20. A method of restoring a packet oriented client signal from at least one synchronous network payload, the method comprising the steps of:

receiving said synchronous payload;

de-mapping said signal from said synchronous payload;

storing said signal in an egress buffer; and

processing said signal to add at least one ordered set provided according to a protocol of said packet orientated client signal, wherein said method of restoring the client signal maintains the integrity of the payload of said packet oriented client signal and preserves a buffer-to-buffer flow control mechanism of said client signal.

21. A method as claimed in claim 20, wherein said step of de-mapping includes removing at least one padding character added to said signal prior to being mapped to said synchronous payload.

22. A method as claimed in claim 20, wherein said at least one ordered set is a client signal idle inserted between client signal packets in said signal according to the client signal protocol.

23. A method as claimed in claim 20, said at least one ordered set is a primitive sequence inserted to form a series of primitive sequences in accordance with the client signal protocol.

24. Apparatus adapted to perform steps in a method of mapping a client signal comprising a packet oriented client signal that uses a buffer-to-buffer flow control mechanism to a synchronous transmission network payload, the apparatus comprising:

a processor for processing said client signal to remove at least one ordered set provided according to a protocol of said client signal to form a second signal;

a buffer for storing the processed client signal in an ingress buffer; and

a mapper for mapping the processed client signal to said synchronous payload,

wherein said apparatus preserves the buffer-to-buffer flow control mechanism of the client signal and maintains the integrity of the payload of the client signal.

25. Apparatus as claimed in claim 24, wherein the apparatus is provided in a network element supporting said client signal.

26. Apparatus as claimed in claim 24, wherein the apparatus is provided in a network element supporting said synchronous network payload.

27. A network element comprising apparatus as claimed in claim 25.

28. A network element comprising apparatus as claimed in claim 26.

29. (cancelled)

30. A method of load balancing traffic comprising a packet orientated client signal across a synchronous network, wherein said traffic comprises at least one synchronous network payload comprising a packet oriented client signal which is controlled by a buffer-to-buffer flow control mechanism, the signal having been mapped to a synchronous network payload, using a method including the steps of: receiving said client signal; processing said client signal to a form suitable for mapping to said payload which preserves a buffer-to-buffer flow control mechanism of the client signal, wherein said step of processing reduces the bandwidth of the client signal while maintaining the integrity of a payload of the client signal; and mapping said processed signal to said synchronous network payload, wherein said method of load balancing comprises the steps of:

pre-allocating an initial bandwidth of said synchronous network payload according to a predetermined condition , wherein said payload comprises a plurality of virtually concatenated virtual containers;

diversely routing said synchronous network payload over said synchronous network; and

in the event of a change in a condition of the network, modifiying the allocated bandwidth.

31. A method of load balancing traffic as claimed in claim 30, wherein bandwidth is automatically modified.

32. A method of load balancing traffic as claimed in claim 30, wherein the bandwidth is automatically modified by the apparatus performing the method of mapping.

33. A method of load balancing traffic as claimed in claim 30, wherein said pre-allocation bandwidth is determined by requirements requested by a user of the network.

34. A method of load balancing traffic as claimed in claim 30, wherein said pre-allocation is automatic.

35. A method of load balancing traffic as claimed in claim 30 wherein said pre-allocation is determined by the condition of the synchronous network.

36. A method of allocating bandwidth in a synchronous digital network for a packet oriented signal having buffer-to-buffer flow control, the method comprising the steps of:

received said packet oriented signal;

processing said packet oriented signal to a processed signal having a form suitable for mapping to a synchronous network payload, wherein the processing preserves a buffer-to-buffer flow control mechanism of said packet oriented signal, wherein said step of processing removes redundant information from the packet oriented signal while maintaining the integrity of a payload of the packet oriented signal;

and mapping said processed signal to a said synchronous network payload having a bandwidth determined according to the bandwidth of said processed signal.